# SIMPLIFIED IT
# THE DELL/BLACKBOARD SOLUTION FOR EDUCATION

A WHITE PAPER SPONSORED BY DELL, INTEL AND BLACKBOARD

BY BLACKBOARD PERFORMANCE ENGINEERING
AND PERFORMANCE TUNING CORPORATION
JUNE 2007

# SIMPLIFIED IT
# THE DELL/BLACKBOARD SOLUTION FOR EDUCATION

## CONTENTS

## EXECUTIVE OVERVIEW

The **Blackboard Academic Suite™** is designed to meet the teaching, learning and outcomes needs of the 21st Century campus. Blackboard is a leading provider of education software for K–12, professional education and higher education.

As more and more administrators, faculty and students began using Blackboard, there was a need for a scalable, reliable, cost-efficient infrastructure. To address this need, the Blackboard Performance Engineering Team has designed a Scalable Reference Architecture.

Blackboard and Dell have worked together to provide a hardware and software deployment solution that implements this Scalable Reference Architecture across multiple nodes. This multi-node implementation features software components from Blackboard, hardware components from Dell and database components from Oracle. Blackboard applications provide clustering and load balancing features for building a scalable, reliable environment. Oracle Database 10g® Real Application Clusters® provide the database scalability and reliability that is critical for meeting the performance and reliability goals of the Scalable Reference Architecture. To complete the Solution, Dell server hardware and Dell/EMC storage systems provide robust systems for Blackboard deployment. In addition, Dell and Blackboard offer consulting services to plan and implement Academic Suite installation.

Recently, a set of tests was conducted to measure the scalability of the Dell/Blackboard Solution. The first set of tests paired single application server nodes with single database nodes. All tests were conducted on the Blackboard Academic Suite Release 7 Application Pack 2, using the Red Hat® Enterprise Linux® AS Operating system. For the initial tests, single-node Oracle databases were utilized, rather than Oracle Real Application Clusters. The Application Servers were Dell PowerEdge™ 1855 blade servers, while the database servers were Dell PowerEdge 6850 rack servers. The goal of these tests was to derive the Performance Archetype Ratio — the optimum ratio of application server instances per Application Server node to each database instance. The PAR was calculated as 3.34 – 4.26 Application instances/Database instance, depending on the exact Intel® Xeon® processor used in the Application server for each test.

A second set of tests was conducted against the complete multi-node implementation. These tests featured thirteen Dell PowerEdge 1855 Application Server nodes, four Dell PowerEdge 6850 servers in an Oracle Real Application Cluster configuration, and a Dell/EMC Clariion® Storage Area Network. As predicted, over 30,000 active user sessions per hour were sustained. This demonstrates the scalability of Blackboard applications with Oracle RAC technology. Furthermore, test results indicate that it should be possible to achieve even higher performance with additional Oracle tuning.

This white paper provides details of the Dell/Blackboard Solution for Education, including scalability and reliability features. Sizing guidelines and performance tuning tips are provided. An analysis of the impact of multi-core Intel Xeon processors is provided, along with details for a Reference Deployment. Educational institutions that require a scalable and highly reliable infrastructure for the Blackboard Academic Suite are encouraged to implement the Dell/Blackboard Solution for Education.

## ABOUT BLACKBOARD PERFORMANCE ENGINEERING

The Blackboard Performance Engineering group partners with clients and vendors to test, validate and report on the efficiency, scalability and resiliency of Blackboard products. Blackboard Performance Engineering applies proven methodologies to generate accurate and reliable data about all the products in the Blackboard Academic Suite, including the Blackboard Learning System — Vista Enterprise License and the Blackboard Learning System — CE Enterprise and Basic License. Blackboard Performance Engineering applies the data from its studies to generate optimal enterprise-class solutions for a variety of academic computing environments.

## BUILDING THE 21ST CENTURY CAMPUS

Blackboard is a leading provider of enterprise-class teaching, learning and outcomes-based software applications and services. Blackboard software is used daily by administrators, faculty and students, as well as digital communities around the world. Blackboard's flagship product, the Blackboard Academic Suite[1], includes the following applications that help institutions transform teaching and learning to reach new audiences:

- The Blackboard Learning System™: A feature-rich and highly customizable Course Management System.[2]
- The Blackboard Community System™: An advanced portal with tools for customizing the presentation of content and administration of data based on user context.[3]
- The Blackboard Outcomes System™: A planning and assessment tool for measuring and improving the success of learning initiatives.[4]
- The Blackboard Content System™: A campus-wide workspace for storing, developing, and displaying content individually and collaboratively.[5]
- The Blackboard Portfolio™ System: A presentation tool for archiving and presenting achievements.

Blackboard has developed a vision of how educators can build the 21st Century Campus[6] after interviewing 50 Higher Education leaders across the United States and Canada. Blackboard's vision of helping schools in *Building the 21st Century Campus* creates opportunities for institutions to improve student engagement and student life, enhance institutional accountability, respond to the demands of a global economy and diversified learning communities and strengthen the management of institutional cost and funding strategies.

Dell has over 20 years of experience developing solutions for colleges and universities all across the country. With their innovative technology and services, they have established a comprehensive framework for improving operational efficiency and educational outcomes.

Technology is helping institutions expand their learning objectives to embrace a greater number of students that are no longer classified as K–12, undergraduate or graduate. Distance learning, global learning partnerships and consortiums, and lifelong learning programs are opening education to new audiences while the traditional educational model is transforming to support the needs and expectations of the 21st century student.

To fulfill the vision of Building the 21st Century Campus, Blackboard Academic Suite enterprise implementations will be required to have certain key characteristics. These characteristics include the following:

- Scalability: Ever-growing user communities and activity levels require a scalable infrastructure. The ability to scale horizontally by adding modular components is highly desirable.
- Reliability: Hardware and software infrastructure should contain no possible single points of failure. Backup and recovery and disaster recovery methods must be defined and implemented.
- Cost-Effectiveness: Hardware should be based upon open standards that are cost-effective to implement, rather than expensive proprietary standards.

Blackboard has designed these characteristics into the architecture of the Blackboard Academic Suite.

## BLACKBOARD SCALABLE REFERENCE ARCHITECTURE

Blackboard has created a Scalable Reference Architecture that is designed to be reliable and cost-effective. There are several different hardware and software components that are necessary to achieve these objectives. Note that some components are mandatory to meet the objectives, such as Oracle Database 10g Real Application Clusters. However, there is freedom to customize some components for individual site requirements, such as the exact server and storage hardware models, as long as the solution matches the needs and anticipated growth of the Blackboard Academic Suite on campus.

1 Blackboard Client Collateral Center, 2004, Reference 11
2 Blackboard Client Collateral Center, 2004, Reference 13
3 Blackboard Client Collateral Center, 2004, Reference 5
4 Yaskin, June, 2006, Reference 2
5 Blackboard Client Collateral Center, 2004, Reference 12
6 Blackboard Client Collateral Center, January, 2007, Reference 6

The Dell/Blackboard Solution for Education applies the Blackboard Reference Architecture to create an optimal hardware and network solution for the Blackboard Academic Suite using Dell hardware servers, Dell/EMC storage devices, and Dell/Brocade® switches. The Dell/Blackboard Solution for Education relies on supporting hardware and software including Oracle Database 10g software and Oracle Database 10g Real Application Clusters, Quest® maintenance and monitoring tools and F5® Big-IP® load balancers.

## OVERVIEW OF THE DELL/BLACKBOARD SOLUTION FOR EDUCATION

To provide an enterprise solution for their common education community customers, Dell and Blackboard have joined as partners to design the Dell/Blackboard Solution for Education. The Dell/Blackboard Solution for Education is a fully realized implementation of the Blackboard Scalable Reference Architecture. The Dell/Blackboard Architecture is illustrated in Figure 1. The following key hardware components are utilized:

- Dell PowerEdge Blade Servers with multi-core Intel Xeon processors for Blackboard applications
- Dell PowerEdge 6850 Servers with multi-core Intel Xeon processors for Oracle Database 10g Real Application Clusters
- Dell PowerConnect Gigabit Ethernet Switches for the Oracle RAC Interconnect
- Dell/EMC Clariion CX3 series Fibre Channel Storage Arrays for database storage
- DellIBrocade Fibre Channel SAN Switches
- Dell PowerVault NX1950 NAS Storage Servers for Blackboard Application file storage
- A monitoring server/SNMP server featuring Quest® Spotlight® for RAC software

Going beyond the Solution design, Dell and Blackboard work together to provide a seamless implementation and support experience for the customer.[7] The implementation features a validated Dell server, SAN hardware and Oracle RAC installations; the same time-tested implementations that Dell has provided for years. The Blackboard applications are installed according to best practices and in accordance with Blackboard Scalable Reference Architecture specifications. Although components from several vendors are included in the Solution, Dell works jointly with Blackboard, Oracle and EMC to provide a user-friendly support path. Users may contact either Dell or Blackboard for a streamlined support experience.

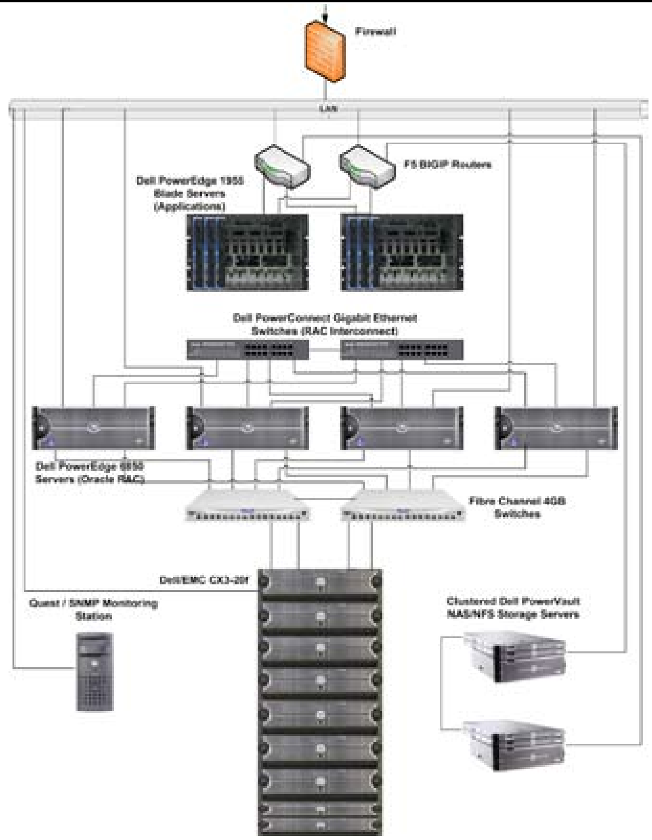


**Figure 1: Dell/Blackboard Architecture**

## SCALABILITY AND RELIABILITY FEATURES OF BLACKBOARD APPLICATIONS

The Blackboard Academic Suite is designed to provide high performance, reliability and scalability. The application architecture is built on enterprise, open source components that include the Apache® Web server and Apache Tomcat® Java® application server. Blackboard applications inherit many of the built-in scalability and reliability features of the Apache architecture, such as distributed workload processing via means of application load balancing and Tomcat application level clustering.

Blackboard provides application level failover capabilities through application level clustering and physical load balancing. If any application instance fails, user requests will be redirected to another instance.

Blackboard application clustering also increases performance by distributing user requests across multiple application instances. Although Java has inherent load balancing capabilities, load

[7] Blackboard Client Collateral Center, 2005, Reference 2

balancing is optimized across servers with the use of one or two F5 Big-IP load balancing routers. These routers maintain virtual IP connections for users to connect to applications and the routers optimize resource utilization by load balancing connections across the pool of virtual IP addresses.

Blackboard also creates multiple Java instances (Java Virtual Machines) per physical server. This greatly enhances scalability by driving resource usage up to optimal levels on each node. Failover and load balancing are maintained for the instances within a node, just as for the instances across cluster nodes.

An enterprise deployment of the Blackboard Application Suite will require multiple application servers. To conserve data center space and resources, 1U–4U rack mounted servers are recommended. Blade server systems are ideal for the application servers because they present optimum processing power in a small amount of space. It is important that each server features fully redundant components, such as dual power supplies, RAID-protected internal disks and so forth.

Multiple, relatively small servers with two to four processor sockets are preferable to large, symmetric, multi-processing (SMP) servers with a large number of processor sockets. There are three reasons for this. First, large SMP servers generally utilize proprietary hardware and operating systems that are complex and expensive in order to scale vertically within the server. Two- to four-socket servers based on open standards are generally much more cost-effective. Second, multiple small servers can effectively scale horizontally through the use of clustering technology. Clustering technology enables scalability and reliability through failover technology. This eliminates single points of failure, which are inherent in the SMP architecture. Lastly, it is easier to upgrade industry-standard building blocks with the latest SMP-compatible micro-processing technology.

## SCALABILITY AND RELIABILITY FEATURES OF DELL HARDWARE

Dell hardware is designed to meet the highest standards of reliability and scalability. Dell is well known for providing enterprise-class servers that are energy efficient and cost effective. Moreover, Dell offers a full spectrum of storage hardware, network hardware and system management software, all designed with built-in reliability and scalability. Dell also works with partners such as Intel, Oracle, EMC, Red Hat and F5 Networks to provide complete solutions for customers. As a result, all hardware required for the Dell/Blackboard Solution for Education is available from Dell.

The following sections list the hardware required for the Dell/Blackboard Solution. In addition, hardware scalability and reliability features are highlighted.

## DELL POWEREDGE™ SERVER HARDWARE

Dell PowerEdge servers using the recommended architecture are designed to host mission-critical applications and databases in data center environments. All Dell PowerEdge servers have key scalability and reliability features in common:

- Hot-plug, redundant power supplies, memory, PCI Express slots, hard drives and cooling fans so that you can replace components without having to make the server unavailable to its users
- State-of-the-art multi-core processors, with 2–4 sockets available, featuring multi-core Intel Xeon processors
- Up to 64GB of high-quality RAM is provided, with advanced high-availability features such as Memory RAID, mirroring and hot-plug DDR2 SDRAM with Error Correcting Code (ECC) memory plus SDDC
- Multiple internal hot-plug SAS or SATA disks are available, with hardware RAID protection options
- Dell PowerEdge servers are available with operating systems pre-installed, including the Red Hat Enterprise Linux AS Operating System.
- Dell provides built-in server management with OpenManage™ software.
- PowerEdge servers are fully tested and validated in clustered environments, including Oracle Real Application Clusters.
- Dual embedded NICs with failover and load balancing support
- Optional redundant Ethernet switches and other I/O technologies
- Optional enhanced management capabilities with DRAC/MC

As mentioned above, the Blackboard Scalable Reference Architecture calls for multiple application servers, preferably in a blade server configuration. The Dell/Blackboard Solution for Education currently utilizes Dell PowerEdge 1955 blade servers for this purpose. One of the advantages of Dell PowerEdge 1955 blade servers is that they conserve valuable data center space while providing high processing capability. Several processors are available for the PowerEdge 1955 servers, including the 5000 and 5100 dual-core Intel Xeon processor series and the 5300 quad-core Intel Xeon processor series. These processors are ideal for Application Server processing.

Up to ten PowerEdge 1955 blade servers are contained in a 7U rack-mount enclosure. The chassis also provides pass-through access to multiple I/O protocols, including Ethernet, Fibre Channel, and InfiniBand. PowerEdge 1955 servers offer several high availability features, including all of the features listed previously.

If a non-blade, rack-mount form factor is desired for application servers, the PowerEdge 1950 offers equivalent processing capability in a 1U chassis. The PowerEdge 2950 server also has similar processing capability, but offers substantially more expansion capabilities in a 2U rack-mount form factor.

Database performance is critical to achieving Blackboard Academic Suite scalability. Dell's premium database server is the PowerEdge 6850, a 4U rack-mount server. This server provides especially strong database performance in Oracle RAC configurations. PowerEdge 6850 servers feature "cutting edge" 7100 series dual-core Intel Xeon processors. Not only do these processors feature some of the fastest clock speeds (up to 3.4 GHz), the large processor cache (up to 16MB L3 cache per dual-core processor) greatly enhances database performance. These servers also feature memory expandability up to 64GB, which is essential for Very Large Databases. PowerEdge 6850 servers offer several high availability features, including all of the features listed previously.

There is an alternative for database servers. PowerEdge 2950 servers are often used for small to moderate size databases. PowerEdge 2950 servers feature similar processors as PowerEdge 1955 servers, and can contain up to 32GB of RAM.

## STORAGE HARDWARE

Because multiple application and database servers must access Blackboard Academic Suite data at the same time, networked storage is required. There are two types of storage that the Blackboard Scalable Reference Architecture is designed to work with: Storage Attached Networks (SAN) and Network Attached Storage (NAS). SAN connects servers to storage through a dedicated network utilizing the Fibre Channel protocol and connected through specialized Fibre Channel switches. SANs have the advantage of very high performance and high reliability, particularly for database applications, at the cost of extra hardware. NAS systems utilize traditional TCP/IP networks and are better suited to file storage than handling the demanding storage requirements such as database activity.

The addition of networked storage increases overall system reliability. SAN and NAS storage arrays are designed with high redundancy, including dual components such as power supplies, dual battery backups and internal failover between components. Even the network switches and cables are redundant, and host-based software can be used for storage network level failover.

There are two types of data that must be accessed within the Blackboard Academic Suite: database content and application content. Accessing database content presents the greatest challenge. Database content is organized relationally, which means that the database server organizes and retrieves the application data internally and a traditional file data structure is not required. In fact, databases communicate with raw block level data, not files. Fortunately, SANs also communicate with raw block data, which make them ideal for database support.

NAS systems may also be used for database support, but there are significant restrictions. NAS systems communicate in terms of file structures. To transmit database block data, format conversion is required for NAS, which adds overhead. Low-end NAS systems or traditional NFS file sharing systems don't guarantee arrival of database blocks, which can lead to database corruption. High-end NAS systems that guarantee delivery of block-level data are required for Oracle databases.

The other type of data that is accessed within the Blackboard Academic Suite is application content. Application content is non-relational, and is accessed in file formats. Any type of NAS is sufficient for this purpose, including traditional NFS mounts. Using a SAN for this purpose is not recommended, due to the need for additional hardware and the need for specialized file systems for global access.

Dell offers multiple options for external storage, ranging from entry-level storage for small businesses to enterprise storage for mission-critical data. Two options are supported for Oracle RAC database storage: Dell PowerVault Direct Attached Storage and Dell/EMC Storage Attached Networks. For Blackboard application data, Dell PowerVault Network Attached Storage is recommended. The following sections will focus on the scalability and reliability features of these external storage options.

### DELL POWERVAULT™ DIRECT ATTACHED STORAGE

For several years, Dell has offered simple, cost-effective external storage in the form of Dell PowerVault Direct Attached SCSI disk arrays. These arrays are controlled by host-based RAID controllers and associated software. Recently, Dell has updated the Dell PowerVault line to work with Serial Attached SCSI (SAS) disks. This has lead to significant performance gains, based on the higher performance potential of the SAS protocol, compared to the traditional parallel SCSI protocol.

In addition, Dell has been steadily improving the scalability and reliability features of PowerVault Direct Attached arrays. The enterprise version of the Dell PowerVault product line is the MD3000 Modular Storage Array. This storage array features read and write performance as fast as any full featured SAN. To enhance scalability, disk sizes as large as 300GB are now offered, with rotational speeds as high as 15K rpm. Up to 45 disks may now be included in an array, delivering a respectable maximum raw capacity of 4.5TB. In terms of reliability, data on disks may be protected with RAID levels 0 (striping only, no protection), 1, 5 and 10. Dual storage controllers provide redundant enclosure management with failover capabilities. Each controller contains 512MB of battery-backed cache (72 hours) to allow data to be saved to disk in case of power failure. There is hardware and software support for maintaining dual paths from two highly available hosts to the storage array.

Despite the attractive features of the PowerVault MD3000 Modular Storage Array, there is one significant limitation. The MD3000 supports a maximum of four SAS connections. For a cluster designed for high availability, two SAS connections must be allocated per server. This means that only two cluster nodes are supported, which limits the scalability of the cluster.

Nevertheless, the PowerVault MD3000 Modular Storage Array has been validated by Dell and Oracle for use with two-node Oracle Real Application Clusters. For Blackboard Academic Suite installations with low to moderate database scalability requirements, the MD3000 array offers a powerful and cost-effective alternative to fibre channel SAN storage.

### DELL/EMC STORAGE ATTACHED NETWORKS

External storage based on Fibre Channel Storage Area Networks is the most commonly deployed storage type for mission-critical data. This is especially true for Oracle RAC deployments, where fibre channel SANs dominate. Through a partnership with EMC, Dell offers the Clariion CX3/CX series fibre channel storage arrays as their enterprise storage solution. Clariion storage arrays may be populated with fibre channel disks for high performance, or inexpensive SATA disks for archival and disk backup, or a combination of both. Clariion CX3 storage arrays may contain up to 480 drives, for a maximum capacity of 239TB. All data communication occurs at 4Gb Fibre Channel speeds, the fastest storage type currently on the market. In addition to high scalability potential for a single server, the Clariion CX series arrays may connect up to 256 servers.

All Clariion arrays include fully redundant components, including redundant battery backups, dual Service Processors with failover capability, and mirrored cache. Clariion arrays present No Single Point of Failure. Data storage reliability may be increased through the use of optional SnapView™ software to assist with point-in-time backup/recovery. In addition, the use of optional MirrorView™ software for remote data replication offers the highest level of system availability.

In most cases, fibre channel switches are required to route data traffic between servers and storage arrays. Dell resells Brocade and McData switches for this purpose. Up to 53 ports per switch are available. For high availability, two switches (dual fabric mode) are recommended to allow for connection failover, along with host-based EMC PowerPath software, to enable active/active connections and failover capability at the same time.

The fibre channel protocol also requires that special host bus adapters (HBAs) be used for connection to a fibre channel fiber optics network. Dell resells QLogic and Emulex host bus adapters to manage server to switch connections. For redundancy, two HBAs per server are recommended, along with EMC® PowerPath® software.

### DELL POWERVAULT NETWORK ATTACHED STORAGE

At a minimum, the Blackboard Academic Suite requires access to an NFS server in order to share Blackboard application files between application servers. Any NFS server (Linux or Microsoft Windows) can be used for this purpose. One option for configuring an NFS server is the "do-it-yourself" approach. All Dell PowerEdge servers are available with Red Hat Enterprise Linux pre-installed. It is relatively easy to configure basic NFS services. However, a simple NFS server represents a potential Single Point of Failure. For a highly available architecture, NFS servers should be clustered. For Red Hat Linux, this requires the Red Hat Cluster Suite (which must be licensed). Although configuring a clustered Linux NFS server is not necessarily difficult, it is clearly not for everyone.

Dell provides a full product line of PowerVault Network Attached Storage with pre-configured NFS. Dell PowerVault Network Attached Storage features Microsoft® Windows® Storage Server 2003 R2, which includes a fully functional NFS server, MSNFS. It is relatively easy to cluster these servers to achieve a High Availability configuration. The entry level servers are the Dell PowerVault Storage server line. These servers are built upon the Dell PowerEdge 830, 2950 and 2900 servers. These servers are designed to share internal disk storage with other servers through NAS protocols, including NFS. They feature 4–10 SAS or SATA internal disks, with up to 3.0TB of raw storage.

The most comprehensive storage server in the Dell PowerVault Networked Attached Storage line is the Dell PowerVault NX1950 Integrated NAS Solution. The NX1950 provides robust support for multiple Networked Storage protocols including NFS, CIFS, iSCSI and fibre channel SAN. In the default configuration, the NX1950 serves as a NAS gateway for a PowerVault MD3000 storage array. It can also serve as a NAS gateway for a CX3 storage array. The NX1950 has been designed for multi-node clustering. Any of Dell's PowerVault Network Attached Storage servers can be effectively used with the Blackboard Academic Suite, but the NX1950 provides the highest levels of scalability and reliability.

## DELL POWERCONNECT™ NETWORK SWITCHES

All access from the outside world to the application servers and the database servers must pass through a network switching layer. The purpose of this layer is to provide hardware controlled load balancing for the application servers, independent of software.

Network connectivity is an important part of the Dell/Blackboard architecture. The Dell/Blackboard Solution for Education calls for at least one load-balancing router (two preferred for redundancy). The F5 Big-IP router (a partner product that Dell resells for the convenience of their customers) was used in the development of this solution.

Dell also offers a full suite of PowerConnect Network Switches. These switches are available in the form of unmanaged switches, managed switches, Web-managed switches, and Power-Over-Ethernet switches. Network speeds of 10/100 Fast Ethernet and Gigabit Ethernet are available.

For the Dell/Blackboard architecture, network switches are required for two purposes: as general purpose LAN switches and for the Oracle RAC private interconnect. The switch type and speed for general purpose LAN switches is determined by local requirements. In general, 100Mb Fast Ethernet meets the minimum requirements for most Local Area Networks. In contrast, Gigabit Ethernet speed is required for Oracle RAC interconnect switches, due to the high throughput requirements of Oracle Cache Fusion. For the Oracle RAC interconnect, dual network connections from each server are recommended for redundancy. It is also recommended to use two switches for interconnect redundancy. To provide failover capabilities for Oracle RAC interconnect connections, native Linux Network Bonding is utilized. If the default options for Linux Bonding are used, two switches from the Web-managed PowerConnect 2700 switch series are sufficient. However, some of the advanced Linux Bonding options require specific switch level support, which may be provided by switches in the PowerConnect 5000 or 6000 series.

## SCALABILITY AND RELIABILITY FEATURES OF ORACLE DATABASE 10G REAL APPLICATION CLUSTERS

Blackboard Application Suite clustering can certainly aid in increasing application scalability and can also enhance reliability. However, any implementation that relies on a single database instance can only support a limited amount of concurrent activity. In addition, a single database instance presents a potential single point of failure for the entire system.

To solve this problem, Oracle offers Oracle Database 10g Real Application Clusters. Oracle RAC technology allows multiple database servers to act as a single large virtual database. Oracle RAC systems can be scaled simply by adding servers to the database cluster, and one or more servers may fail without failure of the virtual database. This effectively allows the Blackboard Academic Suite to reach its true potential in scalability and reliability. Thus, Oracle RAC technology is a key component of the Blackboard Scalable Reference Architecture and the Dell/Blackboard Solution for Education.

## ORACLE REAL APPLICATION CLUSTERS OVERVIEW

Oracle Database 10g Real Application Clusters offer advanced scalability and reliability features that are unmatched. Other database cluster systems offer simple forms of failover in active/passive mode, where only one cluster node is active at any time. Oracle RAC allows full use of all database nodes in active/active mode, and if any node fails, the remaining nodes will take over all work without disruption. As nodes are added to a cluster system, other database cluster systems show only limited performance gains. Oracle Real Application Clusters show significant performance gains as new cluster nodes are added, with work load balanced across cluster nodes. Figure 2 conceptually illustrates failover and load balancing.

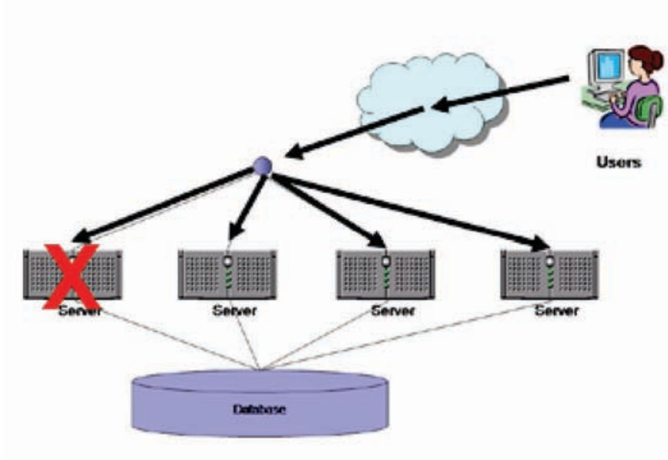


**Figure 2: Oracle RAC Failover and Load Balancing**

In the most succinct terms, Oracle Real Application Clusters consist of 2–10+ database servers with simultaneous read/write access to one set of physical database files. A required component of RAC is storage that is equally shared among all cluster nodes. Shared storage is used to host a single database image that is accessed identically by all cluster servers. Each server includes its own Oracle instance. The definition of an Oracle instance is the combination of runtime processes and dedicated memory that persist throughout the uptime of the database instance.

Of course, it is important that RAC nodes maintain communication with each other, in order to synchronize state information across the cluster. A private interconnect network between cluster nodes allows database servers to share memory structures through a process called Cache Fusion, as conceptually illustrated in Figure 3.

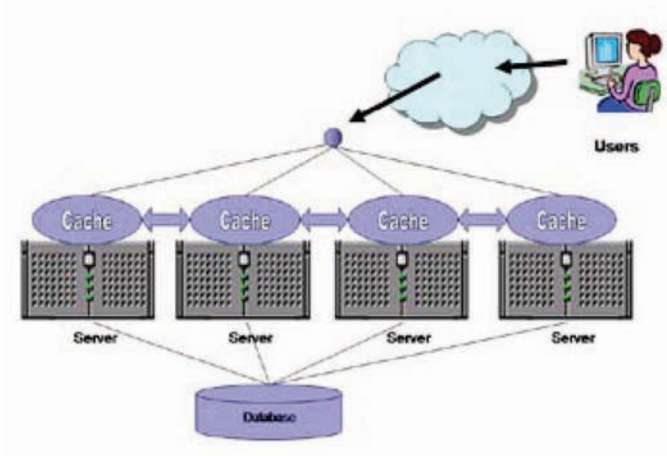


**Figure 3: Oracle RAC Cache Fusion**

Since multiple RAC instances on separate servers can be opened against the same physical database, it is possible to perform work in parallel by simultaneously working within each database instance. Each server performs work on separate user connections. The overall workload is automatically balanced by Oracle Net Services, which assigns user connections to specific servers.

In some cases, it is appropriate to assert manual control over which servers process individual portions of the workload. Server workload can be divided among servers based on rules associated with named User Defined Services (Figure 4). In this example, each node of the four node cluster is associated with a different service. Each service is designed around the requirements for a specific application, such as General Ledger or Human Resources. The users for each service are automatically directed to the appropriate server. Partitioning service resources by application type is an effective method for minimizing contention on the Oracle RAC network interconnect, since cross-node message-passing is minimized in this scenario. Lowered contention on the Oracle RAC interconnect can greatly improve performance. In addition, each server can be optimally tuned for the application that it supports. User Defined Services maintains flexibility. If failover or manual relocation occurs, any service can be relocated to any other node in the cluster.

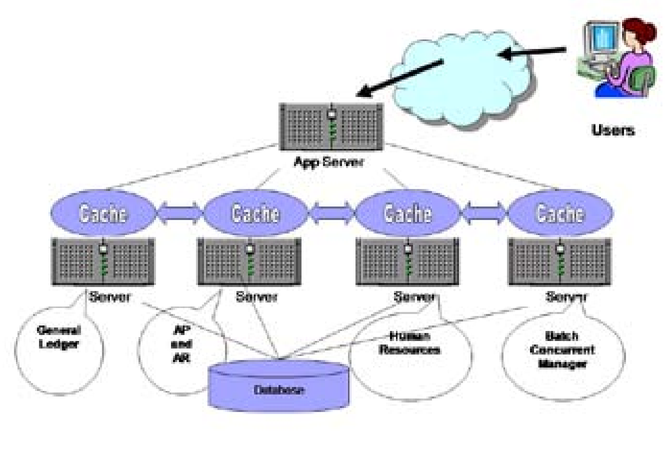


**Figure 4: Partitioning Oracle RAC Workload with User Defined Services**

Whether work is balanced across all servers automatically or User Defined Services are used to control workload, Oracle RAC scales extremely well. Numerous field tests have reported near linear scalability as each new node is added. Figure 5 shows an example from Dell tests conducted with Quest® Benchmark Factory®.

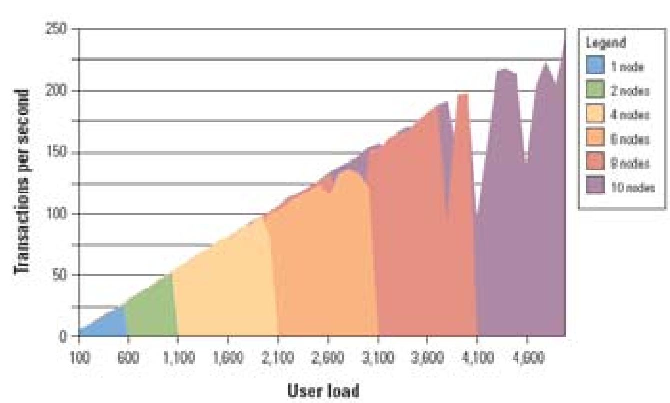


**Figure 5: Quest® Benchmark® Factory Tests of RAC Scalability (Mahmood, et. al, February 2006; Reference 10)**

Oracle RAC also provides high reliability. When a node fails, Oracle log-based recovery is designed to ensure that no committed transactions are lost. Oracle's log-based recovery mechanisms have been extended to handle cases where multiple versions of the same data are held in memory on different nodes, but have not been written to disk (except for the log records). Oracle designed it to ensure that the most recent committed version of a row is recovered, regardless of the existence of multiple versions. Even multiple simultaneous node failures should not cause data loss. In the event of a failure, Oracle relocates and redistributes workload to surviving nodes. Malfunctioning nodes are automatically rebooted, and Oracle automatically attempts to restart all services.

Database server failure events have little impact on database user sessions. Oracle has designed a built-in feature of Oracle RAC and Oracle Net Services called Transparent Application Failover (TAF). If the node a user is connected to should fail mid-session, Transparent Application Failover relocates the user's session to another node with minimal disruption.

In Oracle Database 10g, TAF is expedited through the use of Virtual IP addresses. Users connect to VIP addresses, which are initially specific to single nodes, but which can be relocated to any other node, either through failover or Administrator controlled relocation. Users may notice a temporary slowdown while resources are re-mastered to new nodes, but they will stay connected and work should continue uninterrupted. Only uncommitted transactions at the exact moment of failover are rolled back. Most users will never be aware that a database server has failed.

In addition to the reliability provided by Oracle RAC failover, Oracle RAC is fully compatible with all other Oracle database reliability features. Oracle RAC is designed to utilize built-in Oracle RMAN software for database backup/recovery. Oracle RAC databases may be backed up and recovered from any cluster node. Oracle RAC backup/recovery does not need to be any more complex than single-node Oracle backup/recovery. Optional Oracle Data Guard software may be used for remote database replication. Oracle Data Guard works equally well from RAC to single-node databases, single-node databases to RAC databases, or RAC to RAC databases. Oracle RMAN and Oracle Data Guard are highly recommended for use in conjunction with Oracle RAC databases.

## MONITORING AND MANAGEMENT

There are several useful monitoring and management tools for application activity and database activity. They cover the spectrum from built-in Linux OS monitoring and management tools to Java tools, and include multiple Oracle database tools. Several of these methods were utilized during the Blackboard benchmarking process.

The primary source for hardware and server information is Simple Network Management Protocol data. All hardware can be configured to send Simple Network Management Protocol (SNMP) monitoring data to an SNMP server. SNMP messaging allows hardware and server performance characteristics to be monitored remotely. SNMP monitoring is a key component of the Blackboard Scalable Reference Architecture.

One of the more challenging tasks for an administrator is to monitor and manage Oracle databases. The introduction of Oracle RAC technology makes the task even more challenging. Fortunately, Quest® Software offers tools that greatly simplify Oracle database monitoring and management for database administrators (DBAs). These include Quest Foglight® and Quest Spotlight® on Oracle.

Quest Foglight provides end-to-end database monitoring and management tools for DBAs. Foglight offers 24x7 monitoring as well as alerts, alarms and notification of problems in Oracle databases. Foglight includes the following key components:

- A powerful event-notification engine
- Multi-event correlation capabilities
- A historical repository of key metrics that can be used for performance analysis, trend analysis or capacity planning
- A policy engine for tracking Service Level Agreement compliance
- The ability to handle a heterogeneous mixture of databases, systems and applications

Foglight can be used to monitor the availability and performance of both Oracle databases and the Oracle application infrastructure (Web servers, application servers, the network and the OS) continuously. Foglight offers proactive alert capabilities, so that the monitoring of Oracle databases can be unattended.

Quest Spotlight on Oracle RAC is an industry-standard monitoring tool that is specifically designed to display critical performance metrics within real application clusters (RAC). The homepage view provides a cluster-wide view of all nodes. From the cluster view, database administrators can drill down to details on each cluster node. This gives the DBA the ability to detect node-specific performance issues, as well as issues that affect the entire cluster. Spotlight on Oracle RAC is designed to assist the DBA in discovering performance bottlenecks within their Oracle databases. Spotlight provides graphical displays of both real-time and historical data. Spotlight also includes a granular record and playback function. In addition, Spotlight is unique in that it automatically measures baseline performance and sets its own reporting thresholds. It is then ready to alert DBAs when any anomalous activity occurs. Because Spotlight setup is automatic, it enables busy DBAs to get useful information without investing significant setup and tuning time.

Spotlight on Oracle RAC can detect and report performance issues at a variety of levels, including Top Sessions, Inefficient SQL, Locks, Latches, Wait Events and Disk I/O, helping you identify performance problems quickly. Spotlight on RAC also provides deep integration with Oracle-specific features such as Automatic Storage Management and the Flashback Recovery Area. It is also integrated with Quest Knowledge Xpert® to provide expert insight and advice for resolving RAC performance issues. Quest Foglight and Quest Spotlight on Oracle RAC are recommended components of the Blackboard Scalable Reference Architecture.

## TESTING BLACKBOARD SCALABILITY

The benchmark successfully achieved the active user session count via means of application layer clustering, application layer load-balancing and database clustering with Oracle RAC. Below is a summary of the methodology and process used to achieve this performance objective. The Blackboard Performance Engineering team makes use of an internal sizing process called **Performance Calibration**. This testing process is designed to identify the maximum possible workload hardware architectures can sustain while maintaining acceptable response times.

The process begins with testing an arbitrary workload against a single application server and single database server. The application server is not clustered, nor is the database. It is simply a single instance of the application and database.

Native to the Blackboard load-testing infrastructure is the capability to synthetically simulate user abandonment. Virtual users simulated in the testing environment will attempt to record one or more transactions in the system. If a transaction exceeds a certain time requirement, the user will abandon the system immediately. The ability to simulate user abandonment is a unique capability not found in most performance testing frameworks.

The first test is called the **Abandonment Run**. It is designed to identify three data points: peak of concurrency (POC), average of concurrency (AOC) and level of concurrency (LOC). These data points are workloads of virtual users to be used for proceeding simulations with abandonment disabled.

**Peak of Concurrency: A constant workload for two to five minutes in which the rate of abandonment is equal to the rate of arrival. The workload never can be sustained above this metric of users.**

**Level of Concurrency: A constant workload for the majority of the 60 minute simulation in which user abandonment does not occur. The workload never drops below this metric of users.**

**Average of Concurrency: A derived workload based on the average of the combination of the POC and LOC workloads.**

In the chart listed below, virtual users are presented over the course of roughly 70 minutes. The sample periods are three to four minute blocks. The POC can be identified during the four minute period from (00:06 to 00:10) in which 170 virtual users reached a peak level. The LOC can be identified during the thirteen minute period from (00:50 to 01:03). The AOC is then derived from the average of the POC and LOC. In this case, the AOC is 125 virtual users.

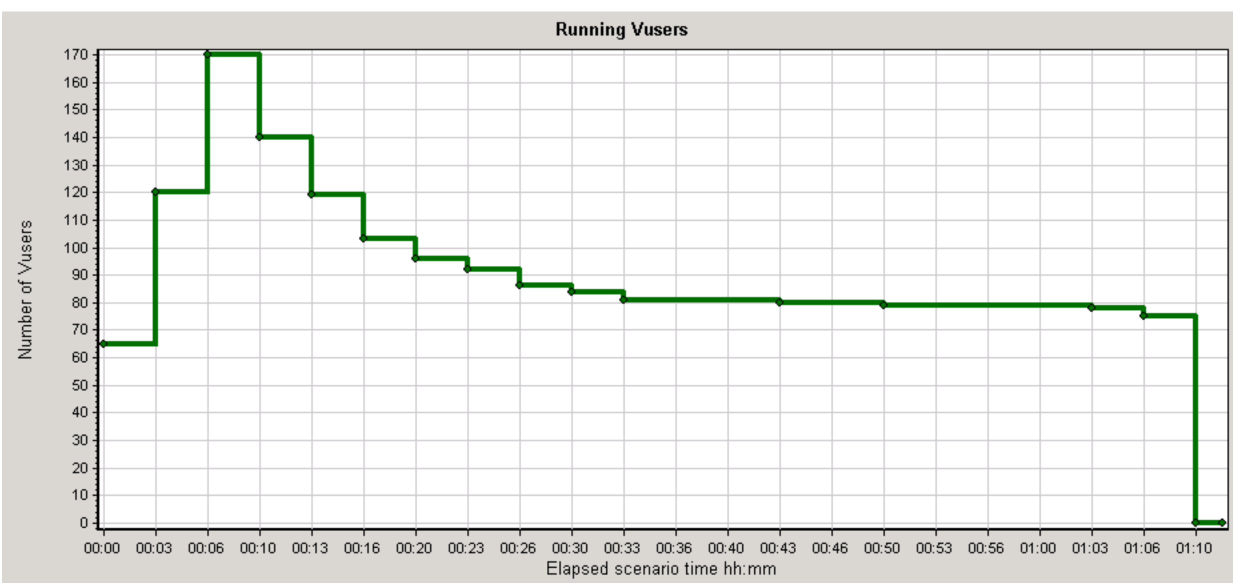


| Color | Scale | Measurement | Graph Minimum | Graph Average | Graph Maximum | Graph Median | Graph Std. Deviation |
|---|---|---|---|---|---|---|---|
| | 1 | Run | 0 | 86.353 | 170 | 84 | 40.708 |

**Figure 6: Example Blackboard Application Suite Abandonment Run Test**

Three subsequent runs are executed with abandonment disabled using the workloads identified as the POC, AOC and LOC. The purpose of these tests is to identify the workload of virtual users that can produce the following:

- Greatest number of active sessions
- Response times below acceptable thresholds
- Application CPU saturation
- Database CPU saturation
- Most consistent hits per second

It is important to note that virtual users do not represent a one-to-one ratio to real users in the application. Virtual users are often an extreme ratio of 15 to 1 or 25 to 1, based on the number of sessions a virtual user can simulate.

This process is to be repeated for every unique piece of application hardware used during the benchmark. In the case of this benchmark, three systems required calibration for achieving the performance objective.

## SINGLE APPLICATION SERVER SIMULATIONS

For the purposes of the benchmark tests, thirteen Dell PowerEdge 1855 Blade Servers were acquired as application servers (this was prior to the release of the PowerEdge 1955 servers). Four Dell PowerEdge 6850 units were acquired as database servers. One of the challenges was a certain amount of heterogeneity in the application servers. Although all of the application servers were PE 1855 blade servers, they featured three different single-core Intel Xeon processors, with processor speeds of 3.6 GHz, 3.73 GHz, and 3.8 GHz. It was decided to proceed with testing with this mix of servers, since it mirrored real world situations where servers are added to a cluster through time, and may not have identical characteristics. In order to perform rigorous testing with the Performance Calibration model, it was necessary to derive three different calibrated application server settings.

For each of the three application server models, the procedure outlined above was performed to arrive at optimum calibration. The metrics tracked were Application Server CPU Utilization, Application Server CPU Standard Deviation, Database Server CPU Utilization, Database Server CPU Standard Deviation and the number of User Sessions per hour. The optimum Application Server CPU Utilization ranged from 75% to 90%, depending on the application server model (Figure 7).

Note that Database Server utilization is always low for the single application server/single database server combination. The data from single server simulations can be used to derive hypothetical performance characteristics of a server architecture assuming performance of the application remains constant. The performance characteristic is called a Performance Archetype Ratio (PAR).

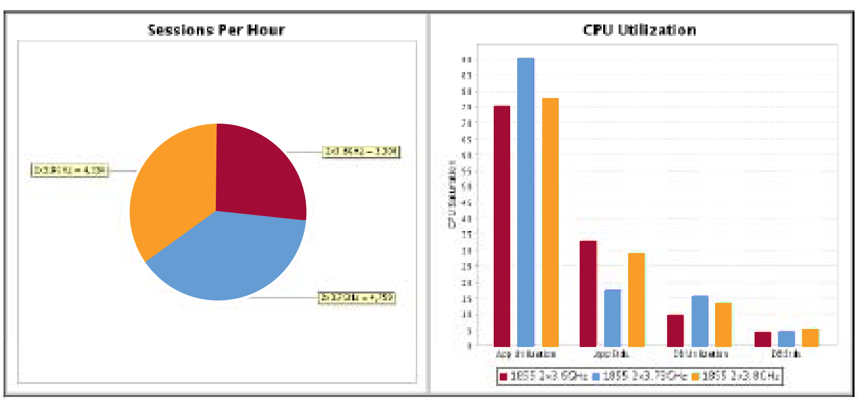


| Server Model | Session Summary | Application Server Utilization | AS CPU SDV | Database Server Utilization | DB CPU SDV | Available Servers |
|---|---|---|---|---|---|---|
| 1855 2x3.6GHz | 3306 | 75.555% | 33.317% | 9.944% | 4.202% | 8 |
| 1855 2x3.73GHz | 4759 | 90.368% | 17.718% | 15.583% | 4.554% | 2 |
| 1855 2x3.8GHz | 4334 | 77.69% | 29.05% | 13.422% | 5.267% | 3 |

**Figure 7: Calibrated Application Server Data for the Dell/Blackboard Single Node Tests**

Performance Archetype Ratios can be called out from any of the metrics identified above. For instance, the Dell 1855 (3.6 GHz) system has a 3,306:1 PAR of sessions to server. The formula for determining the appropriate number of application servers for a desired number of users is simple:

**# APP SERVERS = (# USERS DESIRED) / (# USERS FOR A CALIBRATED APP. SERVER)**

Based on this data point, the number of this particular server model necessary to achieve 30,000 active sessions is roughly 9 to 10 servers. The challenge with this quantification is that the database will become a bottleneck at roughly four application servers. In order to support 9 to 10 application servers, the application environment requires a database infrastructure nearly three times more powerful than the present capacity.

To determine the PAR, it is necessary to determine the optimum Database Server Utilization. In this case, no additional tests were necessary, since the optimum Database Server Utilization (CPU %) is known from many previous tests to be 80%. Given this information, the PAR values may be calculated with the following formula:

**PAR = (OPTIMUM DATABASE SERVER UTILIZATION) / (MEASURED DATABASE SERVER UTILIZATION + 2 STANDARD DEVIATIONS)**

PAR values were calculated to range from 3.24 to 4.36. The following charts in Figure 8 show the PAR data:

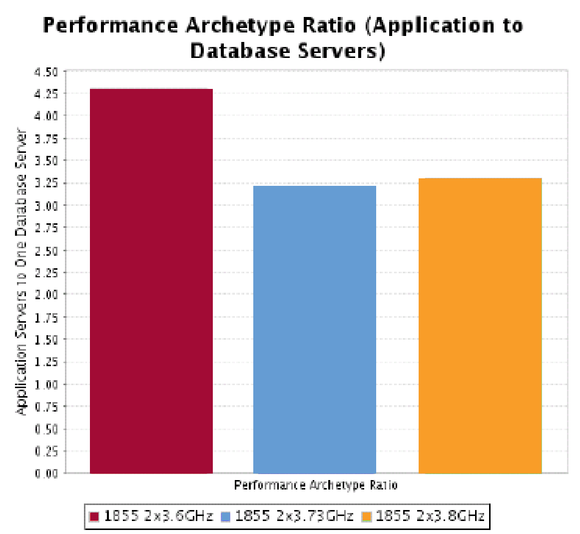


| Server Model | Database Server Utilization | DB CPU SDV x 2 | PAR (APP:DB) |
|---|---|---|---|
| 1855 2x3.6GHz | 9.944% | 8.404% | 4.36 |
| 1855 2x3.73GHz | 15.583% | 10.108% | 3.24 |
| 1855 2x3.8GHz | 13.422% | 10.534% | 3.34 |

**Figure 8: Performance Archetype Ratios for the Dell/Blackboard Single Node Tests**

Based on the Session Summary data values listed in Figure 7, the combination of all thirteen application servers appeared to be more than adequate to achieve the goal of 30,000 sustained sessions, assuming a constant linear scalability.

Assuming all thirteen application servers were to be used, the PAR values were used to estimate the number of database servers required. The PAR value for the Server Model with 3.6 GHz processors was approximately 4:1. Therefore, the eight 3.6 GHz application servers would require approximately two database servers. The 3.8 GHz model had a PAR value of approximately 3:1, and three servers were available, so one database server would be appropriate. Likewise, the model with 3.73 GHz processors also had a PAR of approximately 3:1. Two of these application servers were available, so the fourth database server would be adequate to support these application servers. Thus, a total of four Dell PowerEdge 6850 database servers was required.

It should be noted that the results from these tests were achieved by means of application layer clustering. In order to fully saturate the application servers, multiple JVMs were used on the same physical server.

## MULTIPLE DATABASE NODE TESTS

The final multi-node benchmark test architecture was similar to Figure 1.

- Thirteen Dell PowerEdge 1855 servers were used as application servers.
- Four Dell PowerEdge 6850 servers were used as Oracle Real Application Cluster database servers.
- A Dell/EMC Clariion SAN was used for database storage.
- A Dell PowerEdge server running the Red Hat Linux OS was used as an NFS server for application data.
- An F5 router was used for application server load balancing.
- Dell PowerConnect Gigabit Ethernet switches were used for the Oracle RAC interconnect.

Although the installation process is too detailed to cover here, the following general steps were performed:

1) Place all hardware in racks and provide power.
2) Connect the servers, the SAN, and configure networking.
3) Install the Red Hat Enterprise Linux AS Operating System on each server.
4) Bring up the SAN and check server visibility to the storage.
5) Install the Oracle clusterware.
6) Install the Oracle software.

7) Load/configure the Oracle database.
   a. Three storage options are available:
      i. OCFS2 – Cluster File System, works by default with Blackboard.
      ii. ASM – Automatic Storage Management – requires manual configuration of the Blackboard schema. ASM was used for these tests, according to Dell/Oracle Best Practices.
      iii. NFS – Requires an Oracle certified NAS device (not part of the Dell/Oracle RAC Validated Solution). However, may be used for application file storage.
8) Configure Oracle SID parameters for optimum performance.
9) Test the infrastructure.
10) Load the Blackboard Academic Suite applications.
11) Perform load balancing and failover tests at the application level.

Once full functionality was demonstrated for the entire infrastructure, scalability benchmark tests were conducted. The basic rules were that the test duration was 70 minutes and that peak performance needed to be sustained for at least ten minutes. A successful test was also required to have less than a 1% transaction failure rate and acceptable response times.

As shown in Figure 9, user session counts ranged from approximately 28,000 sessions to approximately 72,000 sessions. To filter the data, outliers were excluded, based on high/low rules and data quality considerations. Tests 4 and 5 were excluded, leaving a range of 27,992 sessions to 38,959 sessions, for an average of 32,437 sessions.

During the benchmark tests, a number of monitoring tools were leveraged in order to understand whether system bottlenecks were limiting application performance. Figure 10 lists and describes the various monitoring tools that were utilized. Most metrics that were monitored showed performance levels that met or exceeded expectations. The various interfaces (application processes, Oracle processes) and resources (JVM memory, Apache server memory size, and Oracle SGA) appeared to be tuned appropriately. One of the concerns going into the benchmark testing was that the Oracle RAC Interconnect might be a performance bottleneck, since passing Cache Fusion memory messages between servers requires high bandwidth. However, during most of the tests the interconnect consistently supported 18 megabytes (2.25 megabits) per second of traffic. From this data, it was concluded that the interconnect did not constrain throughput.

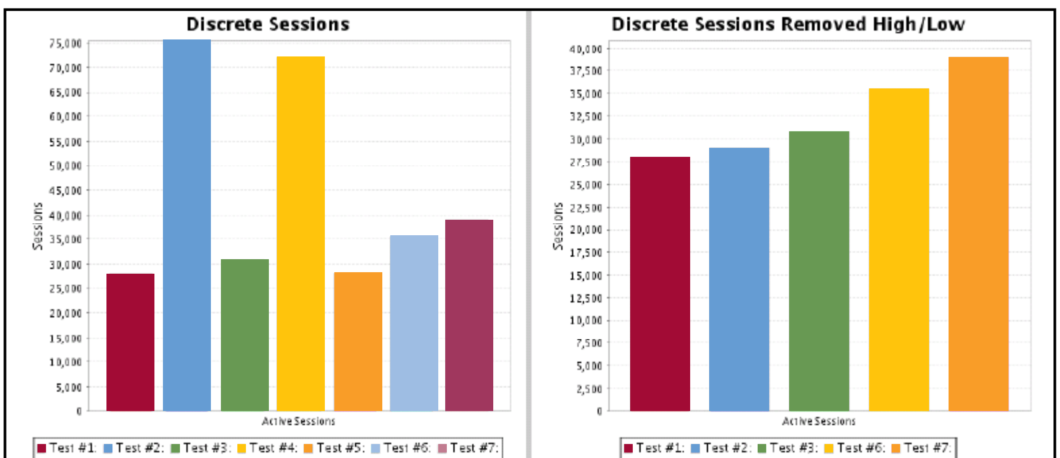


| Test Run | Active Sessions |
|---|---|
| Test #1: | 27992 |
| Test #2: | 28948 |
| Test #3: | 30852 |
| Test #4: | 71989 |
| Test #5: | 28100 |
| Test #6: | 35432 |
| Test #7: | 38959 |

| Test Run | Active Sessions |
|---|---|
| Test #1: | 27992 |
| Test #2: | 28948 |
| Test #3: | 30852 |
| Test #6: | 35432 |
| Test #7: | 38959 |

**Figure 9: Discrete Session Data for Dell/Blackboard Multiple Node Tests**

Nevertheless, the Oracle Statspack and Oracle AWR reports did indicate some specific Oracle RAC object contention issues. There was not sufficient time to revisit these issues during the benchmark tests. However, there is some indication that if these issues are successfully addressed, it may be possible to scale the number of user sessions even higher than the record benchmark results already achieved. The next session focuses on tuning Blackboard performance with Oracle RAC, and some suggestions on how to optimize Oracle RAC performance are provided.

## TUNING BLACKBOARD PERFORMANCE ON ORACLE RAC

Even before the benchmark testing was initiated, it was believed that database performance would be the ultimate control on Benchmark Academic Suite scalability. This belief was reinforced with these benchmark tests. First of all, it is obvious that Oracle RAC allowed scalability significantly beyond what could be achieved with a single database node. It is also well known that Oracle RAC presents new Oracle performance tuning challenges

| TOOL NAME | PURPOSE |
|---|---|
| Apache Server-Status | Apache Server-Status was used to measure counts on active/inactive Apache processes used by the Apache Web Server. Blackboard uses Apache for proxying traffic to an Apache Tomcat Servlet Container, as well as a secondary instance of Apache Web Server for Mod_Perl. |
| JConsole | JConsole is a native tool of the Sun Java SDK library. JConsole was used at periods of instrumentation for determining run-time usage of worker threads and database connections via JDBC. |
| JSTAT | JSTAT is a native tool of the Sun Java SDK library. JSTAT provides instrumentation about memory utilization and pause times (collections) within the various regions of the Java Virtual Machine. |
| -Xloggc | -Xloggc is a lightweight instrumentation tool that provides greater detail of memory consumed by the Java virtualization machine. |
| Top | Top is a native operating system tool available on most Unix systems. Top was used for real-time measurements of CPU utilization and top processes. For Linux systems, Top is an effective tool to watch the performance of KSCAND and KSWAPD. |
| VMstat | VMstat is a lightweight instrumentation tool available on most Unix systems for gathering information about processes, memory, paging, block IO, traps and CPU activity. |
| IOstat | IOstat was used to measure disk performance. |
| Oracle Statspack | Oracle Statspack is a set of SQL, PL/SQL and SQL *Plus scripts which allows the collection, automation, storage and viewing of performance. |
| Oracle AWR | Oracle AWR (Automatic Workload Repository) is a utility much like Oracle Statspack used to collect Oracle performace data. |

**Figure 10: Benchmark Monitoring Tools Used During Blackboard Academic Suite Tests**

and opportunities that are not experienced with single-node Oracle. Without careful performance monitoring and tuning, the full performance potential of Oracle RAC might not be realized.

There are several performance monitoring tools that may be used for tuning Oracle RAC performance:

- Oracle Statspack
- Oracle AWR
- Oracle Grid Control
- Oracle Automatic Database Diagnostic Monitor
- Quest Spotlight on RAC

One tool that was not previously discussed is Oracle Grid Control. Oracle Grid Control is the Oracle Database 10g management framework for Oracle Enterprise Manager (EM) software. Grid Control provides a centralized, integrated approach for managing products in the enterprise. Grid Control contains many self-monitoring features in order to ensure that components are always available and functional.

Every managed server appears as a target on the Grid Control targets page. There is a home page for each server target that provides a consolidated view of resource utilization (CPU, memory and disk). Every managed database also appears as a target on the Grid Control targets page, as well as individual instances.

The other tool that was not previously discussed is Oracle Automatic Database Diagnostic Monitor (ADDM). Oracle ADDM constantly monitors Automatic Workload Repository data and Active Session History data to provide tuning advice. Once every hour, ADDM tuning advice is automatically posted within Oracle Grid Control. Custom ADDM reports can also be created.

One of the potential performance bottlenecks for Oracle RAC is the RAC Interconnect. Even with Gigabit Ethernet speed, interconnect traffic can occasionally exceed the throughput capacity of the RAC interconnect. Interconnect performance can be monitored with the Cluster Interconnects page and the Cluster Cache Coherency page in Oracle Grid Control. If occasional interconnect performance problems occur, it may be possible to use the Oracle tuning tools to trace these problems to specific SQL statements. These statements may then be modified to avoid overloading the interconnect. Another approach is to utilize User Defined Services to connect user sessions to specific database servers. The purpose of User Defined Services is to manage workflow. However, the use of User Defined Services has the side-effect of reducing the amount of inter-node messages, thus reducing interconnect issues.

If interconnect problems are persistent, more drastic solutions are required. It is possible to convert the Linux Bonding interfaces on each cluster node from the default active/passive mode to active/active mode. However, this introduces complexity and may require switch modifications. Dell does not recommend this method. A better long-term solution is to replace the Interconnect network components with InfiniBand components (including switches). The InfiniBand protocol offers inherently faster interconnect speeds.

In the benchmarking tests covered here, a different type of Oracle RAC issue was encountered. Two specific Oracle wait events were repeatedly encountered in the Oracle Statspack and Oracle AWR reports. The two wait events in question were gc buffer busy and enq: TX - row lock contention. Most of the contention appears to be occurring on the BB_BB60.SESSIONS table and the BB_BB60.SESSIONS.SESSION_INDS_PK index. This is due to the fact that every user session frequently performs inserts, updates and deletes against these objects, across multiple RAC nodes. When these types of issues are encountered, there are two generally recommended approaches to mitigating the contention. The first suggestion is to modify the table and index with a hash partitioning scheme. This will spread the table and index over more locations, potentially reducing contention. If the index is still experiencing contention problems, a suggestion is to rebuild the index as a reverse key index. This reverses the byte order of the index leaf node, completely altering the access pattern, and potentially avoiding contention. Addressing these two contention issues can potentially have a noticeable impact on performance.

## REFERENCE DEPLOYMENT

The best thing to take away from the Blackboard benchmark tests is that the Dell/Blackboard architecture provides scalability that would be difficult to achieve with traditional Blackboard implementation methods. There are also some key-sizing best practices that can be gained from these tests. However, it would be a mistake to believe that all Dell PowerEdge servers provide the same processing power as the Application Servers utilized in these tests. In fact, the current generation of Application Servers is already substantially more powerful than the PowerEdge 1855 servers. The following guidelines will address these issues, and should prove useful for those sizing servers for the Blackboard Academic Suite. To assist Blackboard Application Suite customers, system components for a Reference Deployment of the Dell/Blackboard Solution for Education are provided.

## IMPACT OF MULTI-CORE INTEL XEON PROCESSORS

Since the time of the Dell Benchmark, Intel has released both dual-core (5100-series) and quad-core (5300-series) Xeon processors. Recently, the Blackboard Performance Engineering team conducted additional benchmark tests on single application servers. The primary objective of these tests was to quantify the session throughput differences that varying multi-core processors have, compared to single-core processors. Two primary points of data were instrumental: sessions per hour and CPU utilization.

The Blackboard Performance Engineering team analyzed the session throughput performance of eight varying processor models. These models vary from single-core to dual-core to quad-core with differing clock speeds. The test results are illustrated in Figure 11. These results show that dual-core processors are capable of driving up to 6427 active user sessions per hour and quad-core processors are capable of driving up to 6760 active user sessions per hour. In comparison, the best case for single-core processor performance is 4334 active user sessions per hour.

The Blackboard Engineering Team was able to determine some interesting performance trends over the history of the processor releases. The team discovered that the throughput performance of sessions per hour is increasing at a rate of 20% to 30% from model to model. Skipping as many as two generations (2x3.0 GHz single-core to 2x2.33 GHz dual-core) yields a performance factor improvement of nearly 2.5 to 1. This is simply out-of-the-box performance gains with no special optimization or clustering. It should be noted that Intel currently offers several processor models that are significantly more powerful then the processors used in this benchmark. From a dual-core perspective, these models include: 5150 (2.66 GHz) and 5160 (3.0 GHz). From a quad-core perspective, they include: 5320 (1.86 GHz), 5335 (2.0 GHz), 5345 (2.33 GHz) and 5355 (2.66 GHz). The Blackboard Performance Engineering team hypothesizes that similar patterns of performance gains will increase across the varying processor speeds.

The CPU Utilization data serves to underscore the significance of multi-core Intel Xeon processor technology for Blackboard Application Suite performance. Figure 12 illustrates that regardless of clock speed, multi-core technology is able to balance the workload horizontally across sockets and cores, whereas single-core technology with faster clock speeds hits the saturation point faster.

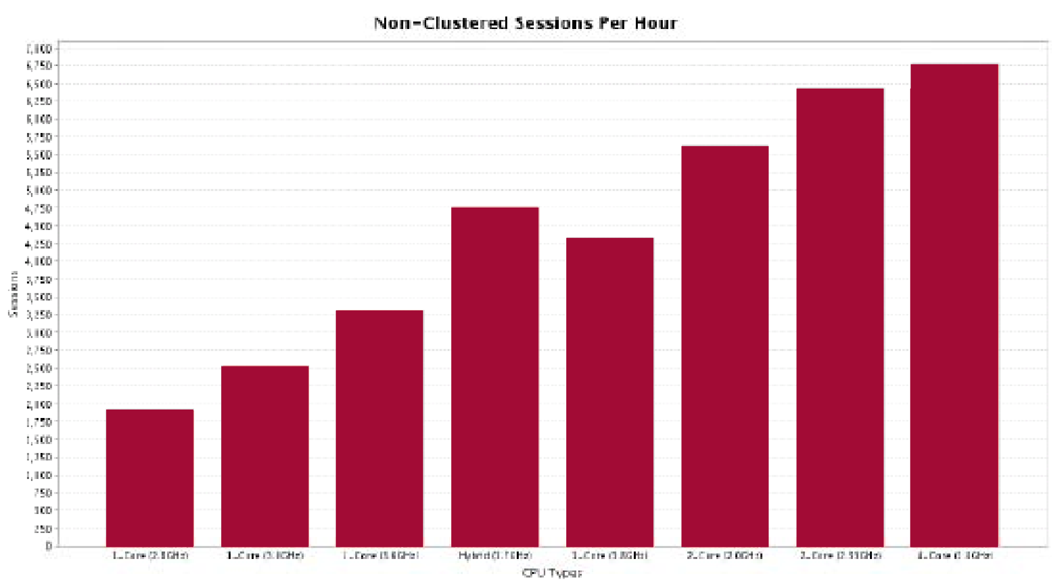


| CPU | 1-Core (2.8GHz) | 1-Core (3.0GHz) | 1-Core (3.6GHz) | Hybrid (3.7GHz) | 1-Core (3.8GHz) | 2-Core (2.0GHz) | 2-Core (2.33GHz) | 4-Core (1.6GHz) |
|---|---|---|---|---|---|---|---|---|
| Active Sessions Per Hour | 1910 | 2520 | 3306 | 4759 | 4334 | 5615 | 6427 | 6760 |

**Figure 11: Comparison of Blackboard Application Suite Throughout Performance for Single-Core and Multi-Core Intel Xeon Processors**

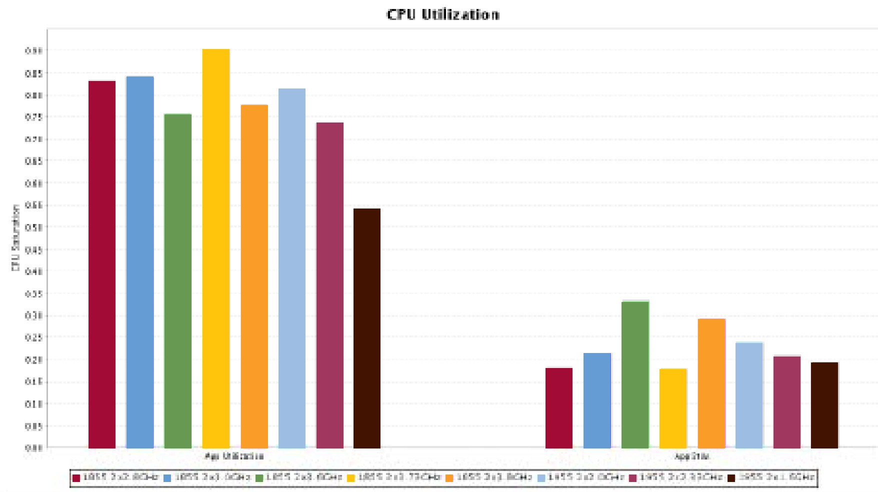


| Server Model | App Utilization | App Stdv. |
|---|---|---|
| 1855 2x2.8GHz | 83.11% | 18.24% |
| 1855 2x3.0GHz | 84.126% | 21.50% |
| 1855 2x3.6GHz | 75.555% | 33.317% |
| 1855 2x3.73GHz | 90.368% | 17.718% |
| 1855 2x3.8GHz | 77.69% | 29.05% |
| 1955 2x2.0GHz | 81.32% | 23.88% |
| 1955 2x2.33GHz | 73.69% | 20.66% |
| 1955 2x1.6GHz | 54.06% | 19.22% |

**Figure 12: Comparison of Blackboard Application Suite CPU Utilization for Single-Core and Multi-Core Intel Xeon Processors**

One of the implications of the relatively low CPU utilization for multi-core processors in the benchmark tests is that there is an opportunity to exploit clustering and virtualization technologies for additional performance gains. Blackboard Performance Engineering has conducted a series of simulations in a clustered configuration, as well as a virtualization architecture using both Xen and VMware ESX edition for Red Hat Linux 4.0. Preliminary results indicate that clustering of Java Virtual Machines provides the greatest session throughput. Further work is in progress in order to quantify the differences of clustering over virtualization.

It is important to note that database performance can control the upper limit of overall system performance for the Blackboard Application Suite. Database servers in the Dell/Blackboard multi-node benchmark tests utilized the dual-core Intel Xeon processors 7120M (3.00 GHz 4MB L3 cache, 800 MHz). Dell PowerEdge 6850 servers are now available with the dual-core Intel Xeon processor 7140M (3.40 GHz 16MB L3 cache, 800 MHz). It is generally known that large level 2 and level 3 processor caches can enhance database performance. Dell has achieved impressive Oracle database benchmark test results with these processors.[8] It is difficult to quantify the effect that these processors will have on database performance for the Blackboard Application Suite. Nonetheless, it is expected that the Dell PowerEdge 6850 servers with the dual-core Intel Xeon processor 7140M will significantly enhance database performance for the Blackboard Application Suite.

## SIZING BEST PRACTICES

One of the sizing concepts that is independent of the server model is resource utilization. It is never a good idea to attempt to achieve 100% resource utilization. In the Blackboard benchmark tests, the optimum Application Server CPU Utilization was 75% to 90%. In general, clients should size all Application Servers to achieve no more than 75% CPU utilization. For database servers, the optimum CPU utilization is 80% in non-RAC mode. In RAC-mode, clients should consider CPU utilization rates around 65% at peak usage periods to allow reserve capacity in case of cluster node failover.

[8] Muirhead, et. al., 2006, Reference 2

One concept not to be overlooked is that hardware sizing should be focused on the anticipated resource utilization at the end of the planning period, not at the beginning. For example, if you expect Blackboard adoption to double within the next two years and hardware is presently being acquired, then the Application and Database servers should be specified for no more than 40% utilization in the first year, so that utilization will be optimized as deployment grows to full capacity.

Sizing I/O resources is a potentially complicated topic. However, there are a few general recommendations for sizing external disk space and I/O:[9]

- The number of disks should be taken into serious consideration. Disk I/O performance is most directly related to the number of disks in a stripe set. Disks should be deployed in groups of ten disks, not individual disks. If some disk space goes unused, it is a reasonable price to pay for good performance.
- Do not drive a single 15K rpm SAS or Fibre Channel disk harder than 125 IO/second (85 IO/second for 10K rpm SAS disks). If you know the IO/second required to support an application, then you can calculate the disks required; i.e. an application that produces 1250 IO/second requires at least ten 15K rpm disks.
- Do not use SATA disks for anything other than archive or backup purposes.
- For the Oracle Cluster File System, create large RAID 10 groups with at least ten groups each. Do not worry about separating Oracle data types.
- For Oracle ASM disks, use multiple small RAID 10 groups, such as 4 disk RAID 10 groups. ASM will create software stripes across the small RAID groups, actually increasing performance.
- In general, do not use RAID 5 for Oracle data. RAID 5 is less expensive to implement, but write performance is generally inadequate for Oracle database performance.

## RECOMMENDED SYSTEM COMPONENTS

There is no single system configuration that will meet the requirements of all Blackboard Academic Suite customers. Nevertheless, the configuration used for the Dell/Blackboard benchmark tests provides a useful reference point for many customers with large user bases, as long as support for 30,000 user sessions per hour is sufficient.

The following list of components has been updated to reflect the latest hardware offered by Dell and the latest multi-core Intel Xeon processors. The system diagram in Figure 1 also reflects these components. It is expected that this configuration should easily support 30,000 active user sessions per hour or more. The recommended components for a Reference Deployment of the Dell/Blackboard Solution for Education are:

- Six Dell PowerEdge 1955 dual-socket blade servers in two chassis, for use as application servers. Quad-core X5355 Intel Xeon processors (2x4MB Cache, 2.66 GHz, 1333MHz FSB) or dual-core 5160 Intel Xeon processors (4MB Cache, 3.0 GHz, 1333MHz FSB) are recommended.
- Four Dell PowerEdge 6850 quad-socket database servers, for use as an Oracle Database 10g Real Application Cluster. Dual-core 7140M Intel Xeon Processors (16MB Cache, 3.4 GHz, 800Mhz FSB) are recommended. As an alternative, Four Dell PowerEdge 2950 dual-socket database servers with quad-core X5355 Intel Xeon processors (2x4MB Cache, 2.66 GHz, 1333MHz FSB) may be utilized.
- A Dell|EMC Clariion Fibre Channel SAN for use as database storage. The CX3-20, CX3-40, or CX3-80 storage arrays are recommended.
- Two Brocade Fibre Channel SAN switches, 16–64 ports each.
- Dual (clustered) PowerVault NX1950 Unified Networked Storage servers for use in storing Blackboard application data.
- Two F5 Big-IP routers for application server load balancing.
- Two Dell PowerConnect Gigabit Ethernet switches for use as the Oracle RAC interconnect. PowerConnect 2700 web-managed Gigabit Ethernet switches are sufficient.
- A Quest/SNMP monitoring station. Any Dell PowerEdge server is sufficient, running Dell OpenManage SNMP monitoring software. Quest Foglight and Quest Spotlight for Oracle RAC are recommended.

[9] Whalen, et. al., 2005; Reference 9

## SUMMARY

Blackboard is currently a leading application provider for Educational Institutions. On each campus, Blackboard software is being utilized by more faculty, staff and students every day. In addition, Blackboard has a clear vision of how Educational Institutions can build the 21st Century Campus.

Blackboard's vision for the 21st Century Campus requires a robust architecture that is scalable, reliable, and cost-effective. Dell has more than two decades of experience developing education solutions and knows the challenges higher education institutions face when implementing new technology.

Together, Dell and Blackboard have created a Scalable Architecture for Blackboard Applications. This architecture contains multiple layers that contribute to the overall level of reliability and scalability. The Blackboard Application layer provides application level clustering and load balancing support. The Oracle Database10g Real Application Cluster layer is designed to provide horizontal scalability and the ability to survive node failure events with little user disruption. Dell provides a hardware layer that features scalable, cost effective PowerEdge Servers and Dell|EMC SAN storage for a robust infrastructure with no Single Point of Failure.

Extensive performance tests were conducted for the Dell/Blackboard infrastructure. The test results indicate that the Dell/Blackboard Solution for Education is highly scalable, supporting up to 30,000 user sessions per hour.

The performance characteristics of multi-core Intel Xeon processors make it possible to use fewer blade servers than were used previously to accomplish a given task, or to scale to higher performance levels than was previously possible. In order to assist customers preparing to move to the Dell/Blackboard architecture, sizing guidelines and a detailed Reference Deployment have been provided. When customers are ready to implement the new architecture, Dell and Blackboard are prepared to offer full installation services and support for the entire Dell/Blackboard Education Solution.

## REFERENCES


1. **"Achieving Continuous Improvement with the Blackboard Outcomes System"**; Author: David Yaskin, Blackboard Inc.; June, 2006, Blackboard Client Collateral Center; http://www.blackboard.com/clientcollateral/Bb%20Outcomes%20System%20-%20Academic%20Bulletin.pdf
2. **"An Enterprise Database Performance Comparison of the Dell PowerEdge 6850 and the Sun Fire V490 UltraSPARC Server"**; Authors: Todd Muirhead and Dave Jaffe, PH.D., Dell Scalable Enterprise Technology Center; October 2006, DellScalable Enterprise Technology Center; http://www.dell.com/downloads/global/power/pe6850_vs_sunv490.pdf
3. **"Blackboard Academic Suite Services"**; Blackboard Client Collateral Center, 2005; http://library.blackboard.com/docs/AS/Bb_Academic_Suite_Datasheet_Services.pdf
4. **"Blackboard Beyond Initiative"**; Blackboard Client Collateral Center, 2007; http://www.blackboard.com/clientcollateral/Beyond_scholar_onesheet.pdf
5. **"Blackboard Community System Capability Set - Detail"**; Blackboard Client Collateral Center, 2004; http://www.blackboard.com/docs/AS/Bb_Community_System_Presentation_Product_Capab_Detail.ppt
6. **"Building the 21st Century Campus"**; Blackboard Client Collateral Center, January, 2007; http://www.blackboard.com/clientcollateral/21st%20Century%20-%20View%20from%20the%20Top%20White%20Paper.pdf
7. **"Intel Performance"**; 2007, Intel Corporation; http://www.intel.com/performance/index.htm
8. **"Oracle RAC The Effectiveness of Scaling Out"**; Author: Erik Peterson, RAC Development Server Technologies, Oracle Corporation; http://www.oracleracsig.org/pls/htmldb/Z?p_url=RAC_SIG.download_my_file?p_file=1000920&p_id=1000920&p_cat=documents&p_user=nobody&p_company=550311706566234
9. **"Oracle Database 10g Linux Administration"**; Author: Edward Whalen, Performance Tuning Corporation, Contributors, Robert Bryan Thomas, Fraser Talbot, Performance Tuning Corporation, Dan Hotka; Oracle Press, Mcgraw-Hill/Osborne; 1st Edition (2005); ISBN 0-07-223053-3
10. **"Testing Oracle 10g RAC Scalability on Dell PowerEdge Servers and Dell/EMC Storage"**; Authors: Zafar Mahmood, Anthony Fernandez, Dell Corporation, Bert Scalzo, PH.D., Quest Software, and Murali Vallath; Dell Power Solutions Magazine, February, 2006; http://www.dell.com/downloads/global/power/ps1q06-20050300-Quest.pdf
11. **"The Blackboard Academic Suite Overview of Product Capabilities"**; Blackboard Client Collateral Center, 2004; http://library.blackboard.com/docs/AS/Bb_Academic_Suite_Whitepaper_Capabilities.pdf
12. **"The Blackboard Content System Overview of Product Capabilities"**; Blackboard Client Collateral Center, 2004; http://library.blackboard.com/docs/AS/Bb_Content_System_Whitepaper_Capabilities.pdf
13. **"The Blackboard Learning System"**; Blackboard Client Collateral Center, 2004; http://library.blackboard.com/docs/AS/Bb_Learning_System_Whitepaper_Capabilities.pdf


Stephen Feldman (sfeldman@blackboard.com) is the Director of Performance Engineering and Architecture at Blackboard Inc. His team is responsible for the performance and scalability of the Blackboard Academic Suite deployed at nearly 3,500 academic institutions world-wide. Stephen constructed Blackboard's Performance Engineering team in the fall of 2003 after serving in a similar capacity with Manugistics, Inc. He has a Bachelors degree in English and Business Administration from Georgetown University. At the time of this paper, Stephen is finishing his Master's degree in Systems Engineering from the Applied Physics Lab at the Johns Hopkins University. Stephen is an active member of the Computer Measurement Group, Academy of Computer Measurement and the Apdex Alliance for Performance Measurement.

Larry Pedigo (lpedigo@perftuning.com) is a Senior Consultant/Instructor for Performance Tuning Corporation. He has been involved with Oracle RAC deployments, SAN deployments, and Oracle and SAN performance optimization for the past three years. Larry has a Bachelors degree in Geology from Texas A&M University, and over 16 years' experience in the IT field.